

Bilinear Games: Polynomial Time Algorithms for Rank Based Subclasses

Jugal Garg^{1*}, Albert Xin Jiang², and Ruta Mehta¹

¹ Indian Institute of Technology, Bombay
jugal,ruta@cse.iitb.ac.in

² University of British Columbia
jiang@cs.ubc.ca

Abstract. Motivated by the sequence form formulation of Koller et al. [19], this paper defines *bilinear games*, and proposes efficient algorithms for its rank based subclasses. *Bilinear games* are two-player non-cooperative single-shot games with compact polytopal strategy sets and two payoff matrices (A, B) such that when (x, y) is the played strategy profile, the payoffs of the players are $x^T A y$ and $x^T B y$ respectively. We show that bilinear games are very general and capture many interesting classes of games like bimatrix games, two player Bayesian games, polymatrix games, two-player extensive form games with perfect recall etc. as special cases, and hence are hard to solve in general.

Existence of a (symmetric) Nash equilibrium for (symmetric) bilinear games follow directly from the known results. For a given bilinear game, we define its *Best Response Polytopes* (BRPs) and characterize the set of Nash equilibria as *fully-labeled* pairs in the BRPs. We consider a rank based hierarchy of bilinear games, where rank of a game (A, B) is defined as $\text{rank}(A + B)$. In this paper, we give polynomial time algorithms to compute Nash equilibrium for special classes of bilinear games:

- Rank-1 games (i.e., $\text{rank}(A + B) = 1$).
- FPTAS for constant rank games (i.e., $\text{rank}(A + B)$ is constant).
- When $\text{rank}(A)$ or $\text{rank}(B)$ is constant. This improves the results by Lipton et al. [22] and Kannan et al. [17], for bimatrix games with low rank matrices.

1 Introduction

In the last decade, there has been much research at the interface of computer science and game theory (see e.g. [24,26]). One fundamental class of computational problems in game theory is the computation of *solution concepts* of finite games. Nash [23] proved that in any finite game there always exists a steady state, from where no player gains by unilateral deviation. Such a steady state has since been named Nash equilibrium (NE) and is perhaps the most well-known and well-studied game-theoretic solution concept. However, computing a Nash equilibrium is nontrivial, and indeed the recent series of papers [4,7,11] established that the problem is PPAD-complete for finite games in the standard *normal form* representation, even for games with only two players. Furthermore it is PPAD-complete to even find a $\frac{1}{n^{\theta(1)}}$ -approximate Nash equilibrium [5]. In light of these negative results, one direction is to identify subclasses of games for which the problem is tractable.

A two-player normal form game can be represented by two payoff matrices, say A and B , one for each player, and hence is also known as bimatrix game [24]. For bimatrix games, polynomial time NE computation algorithms are known for many subclasses, including zero-sum games [6], (quasi-) concave games [20], and games with low rank payoff matrices [22]. A line of work focuses

* Work done while the author was an intern at Microsoft Research India

on games of small *rank*, defined as $\text{rank}(A + B)$ by Kannan and Theobald [17]. They gave a fully polynomial time approximation scheme (FPTAS) for fixed rank games and recently Adsul et al. [1] gave a polynomial time algorithm for computing an exact Nash equilibrium for rank-1 games. Specifying the two payoff matrices of a bimatrix game requires a polynomial number of entries in the numbers of pure strategies available to the players. This is adequate when the set of pure strategies are explicitly given. However, there are situations where the natural description gives the set of pure strategies implicitly, and as a result they may be exponential in the description of the game. For example, normal form (bimatrix) representation of two player extensive-form game may have exponentially many strategies in the size of the extensive-form description [9]. In such a case, even if the resulting bimatrix game has a fixed rank, the above results may not be applied for efficient computation.

Nevertheless, certain types of extensive-form games have some combinatorial structure which can be exploited. Koller, Megiddo and von Stengel [19] converted an arbitrary two-player, perfect-recall, extensive form game into a payoff-equivalent two-player game with continuous strategy sets. In this derived formulation, which they call the *sequence form*, there is a pair of payoff matrices A and B , one for each player. Further, their strategy sets turn out to be compact polytopes in Euclidean space of polynomial dimension. Given a pair of strategies (x, y) , utilities of the players are $x^T Ay$ and $x^T By$ respectively. Interestingly, the sequence form requires only a polynomial number of bits to specify.

Motivated by the sequence form of Koller *et al.*, we define *bilinear games*, which are two-player, non-cooperative, single shot games represented by two payoff matrices, say A and B , of dimension $M \times N$ and two polytopal compact strategy sets $X = \{x \in \mathbb{R}^M \mid Ex = e, x \geq 0\}$ and $Y = \{y \in \mathbb{R}^N \mid Fy = f, y \geq 0\}$. If $(x, y) \in X \times Y$ is the played strategy profile, then $x^T Ay$ and $x^T By$ are the utilities derived by player one and player two respectively. In other words, the payoffs are bilinear functions of strategies, hence the name *bilinear games*. The scope of bilinear games is large enough to capture many interesting classes of games besides two-player extensive form games with perfect recall. For example, for two-player Bayesian games [15,25], polymatrix games [14], and various classes of optimization duels [16], researchers have proposed polynomial-sized payoff-equivalent formulations which (either explicitly or implicitly) turn out to be bilinear games (see Section 2.1 for details). Intuitively the polytopal strategy sets are concise representations of the original sets of mixed strategies as marginal probabilities, and $x^T Ay$ and $x^T By$ express the expected utilities of the original game in terms of these marginal probabilities.

Remark 1. Note that our formulation can express arbitrary polytopes as strategy space, for example if the strategy set of a player is expressed as $\{x : Gx \leq g\}$, *i.e.*, the intersection of a set of half-spaces, it can be transformed to an equivalent game with strategy set $\{x' : Ex' = e, x' \geq 0\}$ using standard techniques (*i.e.*, by adding slack variables, substituting unbounded x_i with $x_i^+ - x_i^-$, $x_i^+, x_i^- \geq 0$, and modifying the payoff matrices accordingly).

As we have seen that many different types of games can be concisely described as bilinear games, designing efficient algorithms for the more general bilinear games seem to be important as well as challenging. Since bimatrix games is a subclass of bilinear games (see Section 2.1), all the hardness results of bimatrix games automatically apply to bilinear games as well. Therefore, the only hope is to design efficient algorithms or FPTAS for the special subclasses. There are many similarities between bilinear and bimatrix games, for example, payoffs are represented by two matrices, and utilities are bilinear functions of strategy vectors, hence it is natural to try to adapt algorithms for bimatrix games to bilinear games. However, a technical challenge is that the polytopal strategy sets of bilinear games are generally much more complex than the bimatrix case; in particular the number

of vertices may be exponential, while the set of mixed strategies is just a simplex. Recently, we were pointed to the constrained games, similar to bilinear games, considered by Charnes [3]. The linear programming technique by Charnes [3] also works for zero-sum bilinear games, *i.e.*, those with $(A + B) = 0$ (see also [18] for sequence form, [16] for another derivation, [25] for zero-sum Bayesian games and [8] for zero-sum polymatrix games). Further, it is easy to show that the linear complementarity program (LCP) characterization for the set of NE of a sequence form game [19] works for bilinear games as well. However the Lemke-Howson algorithm may not be directly applied to the general bilinear games. There are results for certain subclasses of bilinear games with specific structure in their strategy sets, e.g., Howson and Rosenthal [15] adapted the Lemke-Howson algorithm to two-player Bayesian games and Koller *et al.* [19] adapted Lemke’s algorithm [21] to two-player extensive-form games.

Our Contribution. In Section 2, we define the bilinear game and show that the existence of a (symmetric) Nash equilibrium in a (symmetric) bilinear game directly follows from the known results. Note that given a strategy of a player, the other player would like to play a utility maximizing strategy. We formulate this problem as a primal-dual LP, similar to the Koller *et al.* formulation. Using the complementarity conditions of the primal-dual, we characterize the Nash equilibria and then define *Best Response Polytopes* (BRPs) and the notion of *fully-labeled* pairs in BRPs. Further, we show one-to-one correspondence between the Nash equilibria and fully-labeled pairs. This in turn gives a quadratic programming (QP) formulation for the NE computation problem.

Next, we extend Kannan and Theobald’s [17] rank-based hierarchy for bimatrix games to bilinear games, by defining the rank of a bilinear game with payoff matrices (A, B) as the rank of $(A + B)$. Zero-sum games are rank-0 games and a NE for these games can be computed efficiently, as discussed above. In Section 3, we show that in spite of a very general structure of the strategy sets in bilinear games, the basic approach given by Adsul *et al.* [1] to compute a NE of a rank-1 bimatrix game, can be generalized to compute a NE of a rank-1 bilinear game by solving a rank-1 QP. While solving a rank-1 general QP is NP-hard [12], those arising from the bilinear games can be solved in polynomial time. In Section 4, we discuss two FPTAS algorithms for the fixed rank bilinear games, which are generalization of the algorithms by Kannan and Theobald [17] for the bimatrix games. Finally, in Section 5, using the structure of BRPs, we obtain a polynomial time algorithm for the case when the rank of either A or B is a constant, and rank of E and F are also constant. Since, a bimatrix game can be thought of as a bilinear game with E and F being a single row of 1s (see Example 1 of Section 2.1), this algorithm improves upon a result by Lipton *et al.* [22] and Kannan *et al.* [17] for bimatrix games, where they require the rank of both A and B to be constant. This approach also gives an enumeration algorithm for extreme equilibria, which runs in polynomial time under the above assumption and exponential time for the general bilinear games.

The following table summarizes all the NE computation results while keeping the bimatrix games in perspective.

Results	Bimatrix games	Bilinear games
Existence of (symmetric) NE	Nash [23]	Easy to show using ([23]) [10]
LCP formulation for NE	Known [24]	Koller <i>et al.</i> [19]
NE as fully-labeled pairs of BRPs	Known [27]	This paper
Zero-sum games	Linear programming [6]	Linear programming [18]
Rank-1 games	Adsul <i>et al.</i> [1]	This paper
FPTAS for fixed rank games	Kannan and Theobald [17]	This paper
Games with low rank matrices	Lipton <i>et al.</i> [22]	This paper

2 Bilinear Games and Nash Equilibria

Notations. We consider a vector x as a column vector by default and for the row vector, we use transpose (i.e., x^T). A “0” in the block representation of a matrix, is the matrix with all zero entries of appropriate dimension, and “ 1_k ” is a vector of all 1s of length k . Let $x \in \mathbb{R}^n$ be a vector and $c \in \mathbb{R}$ be a scalar, then by $x \leq c$ we mean, $\forall i \leq n, x_i \leq c$. For a given matrix X , X_i denotes the i^{th} row of X , X^j denotes the j^{th} column of X and $|X|$ denotes the maximum absolute entry in X , i.e., $|X| = \max_{ij} |X_{ij}|$. For a set S , $\Delta(S)$ denotes the set of probability distribution vectors over the elements of S , i.e., $\Delta(S) = \{x \in \mathbb{R}^{|S|} \mid x \geq 0, \sum_{i \in S} x_i = 1\}$.

Bilinear games are two-player non-cooperative, single shot games. A bilinear game is represented by two $M \times N$ dimensional payoff matrices A and B , one for each player, and two compact polytopal strategy sets. Let $S_1 = \{1, \dots, M\}$ be the set of rows and $S_2 = \{1, \dots, N\}$ be the set of columns of the matrices. Let $E \in \mathbb{R}^{k_1 \times M}$ and $F \in \mathbb{R}^{k_2 \times N}$ be the matrices, and $e \in \mathbb{R}^{k_1}$ and $f \in \mathbb{R}^{k_2}$ be the vectors. The strategy set of the first-player is $X = \{x \in \mathbb{R}^M \mid Ex = e, x \geq 0\}$ and the second-player is $Y = \{y \in \mathbb{R}^N \mid Fy = f, y \geq 0\}$. Sets X and Y are assumed to be compact. From a strategy profile $(x, y) \in X \times Y$, the payoffs obtained by the first and the second player are $x^T Ay$ and $x^T By$ respectively.

From a Nash equilibrium (NE) strategy profile, no player gains by unilateral deviation. Formally,

Definition 1. A strategy profile $(x, y) \in X \times Y$ is a NE of the game (A, B) iff $x^T Ay \geq x'^T Ay, \forall x' \in X$ and $x^T By \geq x^T By', \forall y' \in Y$.

As a direct corollary of Glicksberg’s [10] result that there always exists a Nash equilibrium in a game whose players’ strategy spaces are convex and compact, and whose utility function for each player i is continuous in all players’ strategies and quasi-concave in i ’s strategy, we have

Proposition 1. Every bilinear game has at least one Nash equilibrium.

A bilinear game is completely represented by a six-tuple (A, B, E, F, e, f) in general. However, for ease of notation we represent it by (A, B) fixing (E, F, e, f) . Given a strategy $y \in Y$ of the second-player, the objective of the first player is to play $x \in X$ such that $x^T(Ay)$ is maximized, i.e., solve the following linear program [19].

$$\begin{array}{ll} \max : x^T(Ay) & \\ \text{s.t.} & Ex = e \\ & x \geq 0 \end{array} \quad \xrightarrow{\text{Dual}} \quad \begin{array}{ll} \min : e^T p & \\ \text{s.t.} & E^T p \geq Ay \end{array} \quad (1)$$

Note that p_i is the dual variable of the equation $E_i x = e_i$ in the above program. At the optimal point (x, p) of (1), we get $x_i > 0 \Rightarrow A_i y = p^T E^i, \forall i \in S_1$ from the complementarity. A similar condition can be obtained for the second-player, given an $x \in X$. At a Nash equilibrium both the conditions are satisfied, and these characterize the NE strategies as follows: A strategy pair $(x, y) \in X \times Y$ is a Nash equilibrium of the game (A, B) iff it satisfies the following conditions.

$$\begin{array}{ll} \exists p \in \mathbb{R}^{k_1} \text{ s.t. } Ay \leq E^T p & \text{and } \forall i \in S_1, x_i > 0 \Rightarrow A_i y = p^T E^i \\ \exists q \in \mathbb{R}^{k_2} \text{ s.t. } x^T B \leq q^T F & \text{and } \forall j \in S_2, y_j > 0 \Rightarrow x^T B^j = q^T F^j \end{array} \quad (2)$$

The above characterization implies that, a player plays a strategy with non-zero probability only if it gives the maximum payoff with respect to (w.r.t.) the opponent’s strategy in *some sense*. Such

strategies are called the *best response* strategies (w.r.t. the opponent's strategy). Using this fact, we define *best response polytopes* (BRPs), similar to the best response polytopes of a bimatrix game [27].

In the following expression, x, y, p and q are vector variables.

$$\begin{aligned} P &= \{ (y, p) \in \mathbb{R}^{N+k_1} \mid A_i y - p^T E^i \leq 0, \forall i \in S_1; \quad y_j \geq 0, \quad \forall j \in S_2; \quad Fy = f \} \\ Q &= \{ (x, q) \in \mathbb{R}^{M+k_2} \mid \quad x_i \geq 0, \quad \forall i \in S_1; \quad x^T B^j - q^T F^j \leq 0, \forall j \in S_2; \quad Ex = e \} \end{aligned} \quad (3)$$

The polytope P in (3) is closely related to the best response strategies of the first-player for any given strategy of the second-player and it is called the *best response* polytope of the first-player. Similarly Q is called the best response polytope of the second-player. Note that, in both the polytopes the first set of inequalities corresponds to the first-player, and the second set corresponds to the second-player. Since $|S_1| = M$ and $|S_2| = N$, let the inequalities be numbered from 1 to M , and $M+1$ to $M+N$ in both the polytopes. Let the *label* $L(v)$ of a point v in the polytope be the set of indices of the tight inequalities at v . If a pair $(v, w) \in P \times Q$ is such that $L(v) \cup L(w) = \{1, \dots, M+N\}$, then it is called a *fully-labeled pair*. The proof of the next lemma follows using (2).

Lemma 1. *A strategy profile (x, y) is a NE of the game (A, B) iff $((y, p), (x, q)) \in P \times Q$ is a fully-labeled pair, for some p and q .*

A game is called non-degenerate if both the polytopes are non-degenerate. Note that a fully-labeled pair of a non-degenerate game has to be a vertex-pair. Lemma 1 implies that a NE strategy profile has to satisfy the following linear complementarity conditions (LCP) over $P \times Q$ [19].

$$\begin{aligned} ((y, p), (x, q)) \in P \times Q \text{ corresponds to a NE} &\Leftrightarrow x^T(Ay - E^T p) = 0 \text{ and} \\ &\quad (x^T B - q^T F)y = 0 \end{aligned} \quad (4)$$

Clearly, $x^T(Ay - E^T p) \leq 0$ and $(x^T B - q^T F)y \leq 0$ over $P \times Q$ and hence $x^T(Ay - E^T p) + (x^T B - q^T F)y \leq 0$. Simplifying the expression using $Ex = e$ and $Fy = f$ we get $x^T(A + B)y - e^T p - f^T q \leq 0$ over $P \times Q$ and equality holds iff (x, y) is a NE (using (4)). This gives the following QP formulation which captures all the NE of game (A, B) at its optimal points.

$$\begin{aligned} \max: & x^T(A + B)y - e^T p - f^T q \\ \text{s.t.} & ((y, p), (x, q)) \in P \times Q \end{aligned} \quad (5)$$

Symmetric Bilinear Games. Nash [23] proved that any symmetric finite game has a symmetric Nash equilibrium. The concept of symmetry can be straightforwardly adapted to the bilinear games: We say a bilinear game is symmetric if $B = A^T$, $E = F$, and $e = f$. A strategy profile (x, y) is symmetric if $x = y$. A straightforward adaptation of Nash's [23] proof yields the following proposition.

Proposition 2. *Any symmetric bilinear game has a symmetric NE.*

Note that for a symmetric game (A, A^T) , strategy sets X and Y are the same. From (2) and (4), it is clear that a symmetric NE $x \in X$ must satisfy $q = p$, $Ax \leq E^T p$, $x^T(Ax - E^T p) = 0$. This gives the following QP formulation to capture all symmetric NE of a symmetric game (A, E, e) .

$$\begin{aligned} \max: & x^T Ax - e^T p \\ \text{s.t.} & Ax \leq E^T p; \quad Ex = e; \quad x \geq 0 \end{aligned} \quad (6)$$

A bilinear game (A, B, E, F, e, f) can be converted to an equivalent symmetric game (A', E', e') , where

$$A' = \begin{bmatrix} 0 & A \\ B^T & 0 \end{bmatrix}, \quad E' = \begin{bmatrix} E & 0 \\ 0 & F \end{bmatrix}, \quad e' = \begin{bmatrix} e \\ f \end{bmatrix}, \text{ with strategy vector } z = \begin{bmatrix} x \\ y \end{bmatrix}$$

It is easy to check that any symmetric NE of the derived game corresponds to a NE of the original game and vice-versa. In the next section, we discuss reductions of different games to polynomial size bilinear games, which do not seem to be possible with the bimatrix games.

2.1 Examples of Bilinear Games

The simplest subclass of bilinear games is the set of two-player normal-form games (bimatrix games).

Example 1. Bimatrix Games

A bimatrix game (A, B) with $A, B \in \mathbb{R}^{M \times N}$ can be straightforwardly transformed to the bilinear game (A, B, E, F, e, f) where $E^T = \mathbf{1}_M, e = 1$ and similarly $F^T = \mathbf{1}_N, f = 1$. \square

Many other interesting classes of finite games may be formulated as bilinear games. In this section, we provide a few examples where the bilinear formulation are exponentially smaller than a direct bimatrix formulation.

Example 2. Two-player Bayesian Games

In a Bayesian game [13], there is a *type set* associated with each player, which is her private information. The nature draws the type for each player from a joint distribution, which is a common knowledge, and each player gets to know only her own type before choosing an action. The final utilities of the game is determined by types of all the players, and hence are uncertain.

Here we consider the two-player case, where T_i s are the type sets and S_i s are the strategy sets. The joint probability distribution is denoted by p_{ts} for the type profile $(t, s) \in T_1 \times T_2$. Let $S = S_1 \times S_2, T = T_1 \times T_2$, $|T_i| = t_i$ and $|S_i| = m_i$. The utilities are the functions of actions and types, *i.e.*, $u_i : S \times T \rightarrow \mathbb{R}$, hence for every type profile they can be represented by the two matrices. For a type profile (t, s) let A^{ts} and B^{ts} denote the respective $m_1 \times m_2$ dimensional payoff matrices. The strategy of a player is to decide her play for each of her type so that her expected payoff is maximized, *i.e.*, $x : T_1 \rightarrow \Delta(S_1)$ for player one and $y : T_2 \rightarrow \Delta(S_2)$ for player two. For a $t \in T_1$, let x^t denote the mixed strategy given type t .

The *induced normal form* of this game is a bimatrix game in which each pure strategy of a player prescribes an action for each of her types. Thus the size of the induced normal form is exponential in the number of types. However, it can be formulated as a polynomial sized bilinear game as follows

$$A = \begin{bmatrix} p_{11}A^{11} & \cdots & p_{1t_2}A^{1t_2} \\ \vdots & \ddots & \vdots \\ p_{t_11}A^{t_11} & \cdots & p_{t_1t_2}A^{t_1t_2} \end{bmatrix}, B = \begin{bmatrix} p_{11}B^{11} & \cdots & p_{1t_2}B^{1t_2} \\ \vdots & \ddots & \vdots \\ p_{t_11}B^{t_11} & \cdots & p_{t_1t_2}B^{t_1t_2} \end{bmatrix}, E = \begin{bmatrix} \mathbf{1}_m^T & 0 & \cdots \\ 0 & \mathbf{1}_m^T & \\ \vdots & & \ddots \end{bmatrix}, F = \begin{bmatrix} \mathbf{1}_n^T & 0 & \cdots \\ 0 & \mathbf{1}_n^T & \\ \vdots & & \ddots \end{bmatrix}$$

and $e = \mathbf{1}_{t_1}, f = \mathbf{1}_{t_2}$. Given mixed strategies $x^1, \dots, x^{t_1}, y^1, \dots, y^{t_2}$ of the Bayesian game, define $x = [x^1, \dots, x^{t_1}]^T$ and $y = [y^1, \dots, y^{t_2}]^T$. Then $x^T A y$ and $x^T B y$ are exactly the expected utilities of the Bayesian game. This transformation is implicit in Howson and Rosenthal's [15] adaptation of Lemke-Howson algorithm to two-player Bayesian games. \square

Example 3. Polymatrix Games [14]

A polymatrix game is an n -player game in which each player's utility is the sum of the utilities resulting from her bilateral interactions with each of the $n - 1$ other players. Let S_i be player i 's set of pure strategies. The game is represented by the payoff matrices $A^{ij} \in \mathbb{R}^{|S_i| \times |S_j|}$ for all pairs of

players (i, j) . Let $x^i \in \Delta(S_i)$ denote a mixed strategy of player i . Given a strategy profile (x^1, \dots, x^n) , the expected utility of player i is,

$$u_i(x^1, \dots, x^n) = \sum_{j \neq i} (x^i)^T A^{ij} x^j$$

We show that any polymatrix game can be transformed to a symmetric bilinear game such that any symmetric NE of the bilinear game corresponds to a NE of the polymatrix game. Our derivation is adapted from Howson's [14] formulation of NE of polymatrix games as solutions of an LCP. Formally, given a polymatrix game, we define the *induced symmetric bilinear game* as (A, A^T, E, E, e, e) , where

$$A = \begin{bmatrix} 0 & A^{12} & \dots & A^{1n} \\ A^{21} & 0 & & A^{2n} \\ \vdots & & \ddots & \\ A^{n1} & A^{n2} & & 0 \end{bmatrix}, \quad E = \begin{bmatrix} 1_{|S_1|}^T & 0 & \dots & 0 \\ 0 & 1_{|S_2|}^T & \dots & 0 \\ \vdots & & \ddots & \\ 0 & 0 & \dots & 1_{|S_n|}^T \end{bmatrix},$$

and $e = 1_n$. The space of strategy vectors is $x = [x^1{}^T \dots x^n{}^T]^T$.

Proposition 3. *Consider a polymatrix game of n players. The strategy (x^1, x^2, \dots, x^n) is an NE of the game if and only if (x, x) is a symmetric NE of its induced bilinear game.*

Proof. The proof is relatively straightforward, by observing that the respective incentive constraints are equivalent. Thus the problem of finding a NE of a polymatrix game reduces to the problem of finding a symmetric NE of a symmetric bilinear game. Note that an asymmetric NE doesn't correspond to a NE of the polymatrix game. \square

Immorlica *et al.* [16] analyzed several classes of games between two optimization algorithms whose objectives are to outperform each other. The space of pure strategies are the possible outputs of the algorithm which are exponential, however the authors were able to formulate some of these games as zero-sum bilinear games (which they call *bilinear duels*). We describe one example from [16].

Example 4. Ranking Duels [16]

Each player chooses a ranking over m elements. Thus the number of pure strategies is exponential in m . Such a ranking can be represented as a $m \times m$ permutation matrix. By the Birkoff-von Neumann theorem, the space of mixed strategies corresponds to the space of $m \times m$ doubly-stochastic matrices with each row and each column sum to 1. This space can be described by the polytope $\{x \in \mathbb{R}^{m^2} \mid x \geq 0; \forall i, \sum_j x_{ij} = 1; \forall j, \sum_i x_{ij} = 1\}$. Viewing x and y as column vectors, the sizes of the corresponding E, e are polynomial in m . Immorlica *et al.* [16] constructed matrices $A \in \mathbb{R}^{m^2 \times m^2}$ such that the players' expected utilities are equal to $x^T A y$ and $-x^T A y$ respectively. \square

Example 5. Two-player Perfect-recall Extensive-form Games

Extensive-form represents a dynamic game as a tree [26], where every pure strategy of a player prescribes a move at each of the player's information sets. As a result the number of pure strategies may be exponential in the size of the extensive-form description. Fortunately, if we assume perfect recall—roughly, that each player remembers all her past decisions and observations—then there always exists a Nash equilibrium in *behavior strategies*, where each player independently chooses a distribution over actions at each of her information sets. Representation of a behavior strategy

requires space linear in the extensive form. However, the expected utilities of the two-player perfect-recall extensive-form games are not bilinear functions of the behavior strategies. Koller *et al.* [19] proposed the *sequence form*, which is a bilinear game formulation for these games. The number of rows and columns of the matrices A and B , in the bilinear form, are the number of feasible sequences of plays of the first and second player respectively. If a play sequence pair (i, j) leads to a leaf node then the ij^{th} entry of A and B are the payoffs of the first and second player at that leaf node, otherwise it is zero*. A strategy x of the first player is such that, $x(\text{root}) = 1$, and if a sequence σ ends at an information node C , then $\sum_{a \in \text{Actions}(C)} x_{\sigma a} - x_{\sigma} = 0$. Similar conditions hold for a strategy y of the second player. Such a strategy may be transformed to a behavior strategy and vice versa. This gives E, F, e and f . Note that the reduction is polynomial sized and $x^T A y$ and $x^T B y$ are exactly the expected payoffs under the corresponding behavior strategies. We refer readers to [19] for more details. \square

Remark 2. Lemke's algorithm on the LCP formulation of bilinear games terminates with a solution (i.e., not at a ray) if the only non-negative solutions x and y to $Ex = 0$ and $Fy = 0$ are $x = 0$ and $y = 0$, and the payoff matrices are non-positive, i.e., $A \leq 0$ and $B \leq 0$. This result directly follows from [19]. Note that games of all the above examples satisfy these requirements, without loss of generality.

The rank of a game (A, B) is defined as $\text{rank}(A + B)$, and we consider the rank based hierarchy of the bilinear games. The set of rank- k games consists of all (A, B) such that $\text{rank}(A + B) \leq k$. Zero-sum games are rank-0 games, the smallest set in the hierarchy. Koller et al. [18] gave an LP formulation for zero-sum bilinear games, derived from two-player extensive form games with perfect recall. However, their formulation works for general bilinear games as well. Beyond rank-0 games, no polynomial time algorithm is known for NE computation (even for the reduction specific formulations). In the next section, we extend the polynomial time solvability of Nash equilibrium for the rank-1 bilinear games.

For all the algorithms that follow, we make the following assumptions (without loss of generality): 1) The entries of A, B, E, F, e and f are integers, since scaling them by a positive number does not change the set of NE. 2) The equalities $Ex = e$ and $Fy = f$ are all linearly independent because even if we discard the dependent equalities, X and Y do not change. 3) The letter \mathcal{L} denotes the bit length of the input game.

3 Rank-1 Games and Polynomial Time Algorithm

The approach used in this section is motivated by the paper [1]. Given a rank-1 game (A, B) , it is easy to find $\alpha \in \mathbb{R}^M$ and $\beta \in \mathbb{R}^N$ such that $A + B = \alpha \cdot \beta^T$, since any two rows of $A + B$ are multiple of each other. In that case $B = -A + \alpha \cdot \beta^T$. Let $G(\alpha) = (A, -A + \alpha \cdot \beta^T)$ be a parametrized game for a fixed $A \in \mathbb{R}^{M \times N}$ and $\beta \in \mathbb{R}^N$. For any game $G(\alpha)$ the BRP of first-player is fixed to $P(\alpha) = P$ (of (3)) since A is fixed. However, the BRP of second-player $Q(\alpha)$ changes with the parameter. Now, consider the following polytope with x, q as vector variables and λ as a scalar variable:

$$Q' = \{ (x, \lambda, q) \in \mathbb{R}^{M+1+k_2} \mid x_i \geq 0, \forall i \in S_1; \ x^T(-A^j) + \lambda\beta_j - q^T F^j \leq 0, \forall j \in S_2; \ Ex = e \} \quad (7)$$

It is easy to see that $Q(\alpha)$ is the projection of $\{(x, \lambda, q) \in Q' \mid \lambda = x^T \alpha\}$ on (x, q) -space. In other words, $Q(\alpha)$ is a section of Q' obtained by hyper-plane $\lambda = x^T \alpha$. Clearly, Q' covers $Q(\alpha)$, $\forall \alpha \in \mathbb{R}^M$. Number the equations of Q' in a similar way as the equations of Q . Let \mathcal{N} be the set of fully-labeled

* Due to chance moves, the entry may correspond to multiple leaf nodes. In that case the entry stores the expected payoff.

pairs of $P \times Q'$, i.e., $\mathcal{N} = \{(v, w) \in P \times Q' \mid L(v) \cup L(w) = \{1, \dots, M + N\}\}$. Using the definition of fully-labeled pairs, it is easy to check that for a given $((y, p), (x, \lambda, q)) \in P \times Q'$,

$$((y, p), (x, \lambda, q)) \in \mathcal{N} \Leftrightarrow x^T(Ay - E^T p) = 0 \text{ and } (x^T(-A) + \lambda\beta^T - q^T F)y = 0 \quad (8)$$

Lemma 2. Let $(v, w) \in \mathcal{N}$, $v = (y, p)$ and $w = (x, \lambda, q)$.

- For all α such that $\lambda = x^T \alpha$, (x, y) is a NE of $G(\alpha)$.
- For every NE (x, y) of a game $G(\alpha)$, there exists a $(v, w) \in \mathcal{N}$, where $\lambda = x^T \alpha$.

Proof. Since (v, w) is fully-labeled it satisfies, $x^T(Ay - E^T p) = 0$ and $(x^T(-A) + \lambda\beta^T - q^T F)y = 0$. Let α be such that $\lambda = x^T \alpha$ then we get $(x^T(-A) + (x^T \alpha)\beta^T - q^T F)y = 0 \Rightarrow (x^T(-A + \alpha\beta^T) - q^T F)y = 0$. This implies that (x, y) is a NE of the game $(A, -A + \alpha\beta^T)$ (i.e., $G(\alpha)$), since it satisfies the complementarity condition of (4).

Given a (x, y) of $G(\alpha)$, from (2) and (4) it is clear that $\exists p, q$ such that $x^T(Ay - E^T p) = 0$ and $(x^T(-A + \alpha\beta^T) - q^T F)y = 0$. Let $\lambda = x^T \alpha$, then we get $(x^T(-A) + \lambda\beta^T - q^T F)y = 0$. Therefore, $((y, p), (x, \lambda, q)) \in \mathcal{N}$. \square

The above lemma establishes strong relation between the set of NE of all the $G(\alpha)$ s and the set \mathcal{N} . Next we discuss the structure of \mathcal{N} , and later use it to design a polynomial time algorithm to find a NE of a given game $G(\alpha)$.

The polytopes P and Q' are assumed to be non-degenerate³, and let $k_1 = k_2 = k$ for simplicity. As there are k linearly independent equalities in P and Q' , they are of dimension N and $M + 1$ respectively. Therefore, $\forall (v, w) \in \mathcal{N}$, $|L(v)| \leq N$ and $|L(w)| \leq M + 1$. Since, $M + N$ labels are required for a pair (v, w) to be part of \mathcal{N} , $\mathcal{N} \subset 1$ -skeleton of $P \times Q'$. Further, if $(v, w) \in \mathcal{N}$ is a vertex pair then $|L(v) \cap L(w)| = 1$. Let the label in the intersection be called the *duplicate label* of (v, w) . Relaxing the inequality corresponding to the duplicate label at (v, w) in P and Q' respectively gives its two adjacent edges in \mathcal{N} . Therefore, every vertex of \mathcal{N} has degree two. This implies that \mathcal{N} is a set of cycles and infinite paths (unbounded edges at both the ends). We will show that \mathcal{N} forms a single infinite path. The next lemma follows directly from the definition of P (3) and Q' (7), and expression (8).

Lemma 3. For all $(v, w) = ((y, p), (x, \lambda, q)) \in P \times Q'$, we have $\lambda(\beta^T y) - e^T p - f^T q \leq 0$, and the equality holds iff $(v, w) \in \mathcal{N}$.

Lemma 3 implies that \mathcal{N} is captured by $\lambda(\beta^T y) - e^T p - f^T q = 0$ over $P \times Q'$. Using this fact and Lemma 3, we define the following parametrized LP.

$$LP(\delta) = \max: \delta(\beta^T y) - e^T p - f^T q \\ \text{s.t. } P \times Q'; \lambda = \delta$$

For an $a \in \mathbb{R}$, let $OPT(a)$ be the set of optimal solutions of $LP(a)$ and $\mathcal{N}(a)$ be the set of points of \mathcal{N} with $\lambda = a$, i.e., $\mathcal{N} = \{(v, w) \in \mathcal{N} \mid w = (x, \lambda, q) \text{ and } \lambda = a\}$.

Lemma 4. For an $a \in \mathbb{R}$, $\mathcal{N}(a) \neq \emptyset$ and $OPT(a) = \mathcal{N}(a)$

Proof. Consider a game $G(\alpha)$ where $\alpha = [a, \dots, a]$. Clearly, for any Nash equilibrium (x, y) of $G(\alpha)$ the corresponding point in \mathcal{N} has $\lambda = x^T \alpha = a$ (Lemma 2). Therefore, $\mathcal{N}(a) \neq \emptyset$. The feasible set of $LP(a)$ is all points of $P \times Q'$ with $\lambda = a$. Further, function $\lambda(\beta^T y) - e^T p - f^T q$ achieves maximum only at points on \mathcal{N} (Lemma 3). Therefore, $OPT(a) = \mathcal{N}(a)$ follows. \square

³ Degeneracy may be handled using standard techniques as done in [1].

Lemma 5. *The set \mathcal{N} forms an infinite path, with λ being monotonic on it.*

Proof. To the contrary suppose there are cycles and multiple paths in \mathcal{N} . Let \mathcal{C} be a cycle in \mathcal{N} . It is easy to see that $\mathcal{N}(a) = \text{intersection of } \mathcal{N} \text{ with the hyper-plane } \lambda = a$. Therefore, $\exists a \in \mathbb{R}$, such that either \mathcal{C} is contained in $\lambda = a$ or it cuts the cycle at exactly two points. This contradicts that $\mathcal{N}(a)$ is a convex set in both the cases (Lemma 4).

Now let \mathcal{P}_1 and \mathcal{P}_2 be two paths in \mathcal{N} . Since, $\mathcal{N}(a)$ is a convex set $\forall a \in \mathbb{R}$, λ is monotonic on both the paths. Suppose, the range of λ covered by \mathcal{P}_1 and \mathcal{P}_2 be $(-\infty, a]$ and (a, \inf) . However, this contradicts the fact that \mathcal{P}_2 is a closed set. Monotonicity of λ follows from the convexity of $\mathcal{N}(a)$. \square

3.1 Algorithm

Let (A, B) be a given rank-1 game and $A + B = \gamma \cdot \beta^T$. Let $\gamma_{\min} = \min_{x \in X} \sum_{i \in S_1} \gamma_i x_i$ and $\gamma_{\max} = \max_{x \in X} \sum_{i \in S_1} \gamma_i x_i$. The γ_{\min} and γ_{\max} exists since X is a bounded polytope. From Lemma 2 it is clear that every point in the intersection of the set \mathcal{N} and hyper-plane $H_\gamma : \lambda - \sum_{i \in S_1} \gamma_i x_i = 0$ corresponds to a NE of the given game (A, B) . Note that for any point in the intersection, corresponding λ is between γ_{\min} and γ_{\max} . Let H_γ^- and H_γ^+ be the negative and positive half spaces of the hyper-plane H_γ respectively, then clearly $\mathcal{N}(\gamma_{\min}) \in H_\gamma^-$ and $\mathcal{N}(\gamma_{\max}) \in H_\gamma^+$. All the points in the intersection of \mathcal{N} and H_γ are between $\mathcal{N}(\gamma_{\min})$ and $\mathcal{N}(\gamma_{\max})$. The following algorithm does binary search on \mathcal{N} between $\mathcal{N}(\gamma_{\min})$ and $\mathcal{N}(\gamma_{\max})$ to find a point in the intersection using the fact that λ monotonically increases (similar to the algorithm in [1]).

- S₁ Initialize $a_1 = \gamma_{\min}$ and $a_2 = \gamma_{\max}$.
- S₂ If the edge containing $\mathcal{N}(a_1)$ or $\mathcal{N}(a_2)$ intersects H_γ , then output the intersection and exit.
- S₃ Let $a = \frac{a_1 + a_2}{2}$. Let $\overline{u}, \overline{v}$ be the edge containing $\mathcal{N}(a)$.
- S₄ If $\overline{u}, \overline{v}$ intersects H_γ , then output the intersection and exit.
- S₅ Else if $\overline{u}, \overline{v} \in H_\gamma^-$, then set $a_1 = a$ else set $a_2 = a$ and continue from step S₃.

Correctness. Since the feasible set of $LP(a)$ is a section of $P \times Q'$, where $\lambda = a$, the $OPT(a)$ is on an edge of $P \times Q'$ (assuming non-degeneracy of $LP(a)$). It is easy to construct this edge from the tight equations of $P \times Q'$ at $OPT(a)$. Clearly, this entire edge should be part of the set \mathcal{N} , hence if this edge intersects the hyper-plane H_γ then we get a Nash equilibrium of the given game. Since, all the points in the intersection of \mathcal{N} and H_γ are between $\mathcal{N}(\gamma_{\min})$ and $\mathcal{N}(\gamma_{\max})$, and λ monotonically increases between these two (Lemma 5), the algorithm does a simple binary search between γ_{\min} and γ_{\max} to find an a , such that the edge containing $\mathcal{N}(a) = OPT(a)$ intersects H_γ (Lemma 4).

Time Complexity. Recall that \mathcal{L} is the bit length of the input game. Since, γ_{\min} and γ_{\max} are optimal points of two LPs on set $X = \{Ex = e, x \geq 0\}$, they can be represented in $\text{poly}(\mathcal{L}, M, N)$ bits. Let $Z = \max\{|A|, |E|, |F|, |e|, |f|, |\gamma|, |\beta|\}$, $l = M + N + k_1 + k_2 + 1$, and $\Delta = l!Z^l$.

Theorem 1. *The above algorithm finds a NE of game (A, B) in time $\text{poly}(\mathcal{L}, M, N)$.*

Proof. One round of steps S₃ to S₅ can be done in polynomial time since computation of $\mathcal{N}(a)$ requires solving $LP(a)$ (Lemma 4), and computation of $\overline{u}, \overline{v} \cap H_\gamma$ requires checking the feasibility of a polytope. Now, to show polynomial time complexity, we need to bound the number of rounds of steps S₃ to S₅.

Note that the denominator of any co-ordinate of a vertex of $P \times Q'$ is at most Δ , and if λ is not constant on an edge of \mathcal{N} , then the difference in its value between the two end points of the edge is

at least $\frac{1}{\Delta^2}$. Therefore, if $a_2 - a_1 < \frac{1}{\Delta^2}$ the algorithm terminates. After k rounds $a_2 - a_1 = \frac{\gamma_{\max} - \gamma_{\min}}{2^k}$. In round k if $a_2 - a_1 = \frac{\gamma_{\max} - \gamma_{\min}}{2^k} > \frac{1}{\Delta^2}$, then $k < \log(\gamma_{\max} - \gamma_{\min}) + 2 \log \Delta$. Therefore, the algorithm is guaranteed to terminate after $\log(\gamma_{\max} - \gamma_{\min}) + 2 \log \Delta + 1 = \text{poly}(\mathcal{L}, M, N)$ many rounds. \square

4 FPTAS for Rank- k Games

In this section, we discuss fully polynomial time approximation schemes for fixed rank games (i.e., $\text{rank}(A + B)$ is constant). The approximation notion in bilinear games can be defined in a similar way to that of bimatrix games given by Kannan et al. [17]. Let $x_{\max} = \max_{x \in X} \sum_i x_i$, $y_{\max} = \max_{y \in Y} \sum_j y_j$ and $D = |A + B|$. Clearly the total payoff derived from a strategy profile $(x, y) \in X \times Y$ is at most $x_{\max} D y_{\max}$. Using this we define an ϵ -approximate NE for a bilinear game (A, B) as follows.

Definition 2. For a strategy profile $(x, y) \in X \times Y$, let $u = \max_{x' \in X} x'^T A y$ and $v = \max_{y' \in Y} x^T B y'$. Then (x, y) is an ϵ -approximate NE of the game (A, B) if $u + v - x^T(A + B)y \leq \epsilon(x_{\max} D y_{\max})$.

For a bimatrix game $x_{\max} D y_{\max} = D$, since x and y are probability distributions, which is compatible with the definition of [17]. Next we define a stronger notion of ϵ -approximate NE called *relative ϵ -approximate NE*, where the error is relative to the maximum achievable payoff from the given strategy.

Definition 3. For a strategy profile $(x, y) \in X \times Y$, let $u = \max_{x' \in X} x'^T A y$ and $v = \max_{y' \in Y} x^T B y'$. Then (x, y) is a *relative ϵ -approximate NE* of the game (A, B) if $u + v - x^T(A + B)y \leq \epsilon(u + v)$, i.e., the total error is relatively small.

Since the value of $u + v$ is at most $x_{\max} D y_{\max}$, if (x, y) is relative ϵ -approximate NE, then it is also ϵ -approximate NE. For all the examples mentioned in Section 2.1, a (relative) approximate NE of the bilinear game formulation can be straightforwardly turned into a (relative) approximate NE of the corresponding finite game under standard definitions. Without loss of generality we assume that A, B, E, F, e and f are integer matrices, since scaling them by a positive value does not change the set of (relative) ϵ -approximate NE. Next we discuss two FPTAS to solve QP of (5), one for each definition of approximation. The approaches used in these algorithms are generalization of the ones in [17].

4.1 FPTAS for Approximate NE

We show that the result by Vavasis [28] can be applied to get an ϵ -approximate Nash equilibrium (Definition 2). The following proposition states the result by Vavasis.

Proposition 4. Let $\min\{\frac{1}{2}x^T Q x + q^T x : Ax \leq b\}$ be a quadratic optimization problem with compact polytope $\{x \in \mathbb{R}^n : Ax \leq b\}$, and let the rank of Q be a fixed constant. If x^* and $x^\#$ denote points minimizing and maximizing the objective function $f(x) = \frac{1}{2}x^T Q x + q^T x$ in the feasible region, respectively, then one can find in time $\text{poly}(\mathcal{L}, \frac{1}{\epsilon})$ a point x^\diamond satisfying

$$f(x^\diamond) - f(x^*) \leq \epsilon(f(x^\#) - f(x^*)).$$

Now consider the following QP formulation of (5), which captures all the NE of (A, B) at its optimal.

$$\begin{aligned} \min: & e^T p + f^T q - x^T(A + B)y \\ \text{s.t. } & Ay - E^T p \leq 0; \quad Fy = f; \quad y \geq 0 \\ & x^T B - q^T F \leq 0; \quad Ex = e; \quad x \geq 0 \end{aligned}$$

Theorem 2. Let (A, B) be a rank- k game, then for every $\epsilon > 0$, an ϵ -approximate Nash equilibrium can be computed in time $\text{poly}(\mathcal{L}, \frac{1}{\epsilon})$, where \mathcal{L} is the bit length of the game and k is a constant.

Proof. The objective function of the above QP can be easily transformed to the standard QP form $\frac{1}{2}x^T Qx + q^T x$, where $\text{rank}(Q) = 2k$. To apply Proposition 4 on this QP, we need to bound its feasible set. Since, $\{x : Ex = e, x \geq 0\}$ and $\{y : Fy = f, y \geq 0\}$ are compact, the only variables to bound are ps and qs . Since, the maximum possible value of $x^T(A + B)y$ for any $(x, y) \in X \times Y$ is $x_{\max} Dy_{\max}$, the value of $e^T p + f^T q$ is at most $x_{\max} Dy_{\max}$ at any point of the polytope corresponding to NE (by (4)). Therefore, we impose $e^T p + f^T q \leq x_{\max} Dy_{\max}$. However, this may not bound the ps and qs .

Let $Z = \max\{|A|, |B|, |E|, |F|, |e|, |f|\}$ and $l = M + N + k_1 + k_2$. Recall that NE of a non-degenerate game correspond to vertices of the polytope. It is easy to see that maximum absolute value of a coordinate of any vertex in the polytope is at most $l!Z^l$. Further, the quantity $l!Z^l$ can be represented in $\text{poly}(\mathcal{L})$ bits. Therefore, imposing $-l!Z^l \leq p \leq l!Z^l$ and $-l!Z^l \leq q \leq l!Z^l$ in the above QP incur only a polynomial increase in its representation and does not change its optimal set. The minimum and the maximum objective values of this QP are zero (Lemma 3) and at most $2x_{\max} Dy_{\max}$ respectively. Let $((y^\diamond, p^\diamond), (x^\diamond, q^\diamond))$ be the solution given by Vavasis algorithm for $\frac{\epsilon}{2}$, then from Proposition 4 we get,

$$e^T p^\diamond + f^T q^\diamond - x^{\diamond T} (A + B) y^\diamond \leq \epsilon (x_{\max} Dy_{\max})$$

From the primal-dual formulation of (1) it is clear that $\max_{x' \in X} x'^T A y^\diamond \leq e^T p^\diamond$ and $\max_{y' \in Y} x^{\diamond T} B y' \leq f^T q^\diamond$. Therefore, we get $\max_{x' \in X} x'^T A y^\diamond + \max_{y' \in Y} x^{\diamond T} B y' - x^{\diamond T} (A + B) y^\diamond \leq \epsilon (x_{\max} Dy_{\max})$. \square

4.2 FPTAS for Relative Approximate NE

Let the rank of a game (A, B) be k , then $A + B = \sum_{i=1}^k \alpha(i) \beta(i)^T$, where $\forall i, \alpha(i) \in \mathbb{R}^M$ and $\beta(i) \in \mathbb{R}^N$. We assume that the game is such that $\alpha(i)$ s and $\beta(i)$ s are positive vectors. For all $i \leq k$, let $w_i = \min_{x \in X} x^T \alpha(i)$ and $w'_i = \max_{x \in X} x^T \alpha(i)$, similarly let $z_i = \min_{y \in Y} \beta(i)^T y$ and $z'_i = \max_{y \in Y} \beta(i)^T y$. Note that w_i, w'_i, z_i and z'_i can be represented by $\text{poly}(\mathcal{L}, M, N)$ bits, since X and Y are compact. Given an $\epsilon > 0$, consider the sub-intervals $[w_i, (1 + \epsilon)w_i]$, $[(1 + \epsilon)w_i, (1 + \epsilon)^2 w_i]$ of $[w_i, w'_i]$ and similarly of $[z_i, z'_i]$. All combinations of these intervals form a grid in $2k$ -dimensional box $\mathcal{B} = \times_i [w_i, w'_i] \times_i [z_i, z'_i]$. Let $(x, y) \in X \times Y$ be such that $\forall i, x^T \alpha(i) \in [u_i, (1 + \epsilon)u_i]$ and $\beta(i)^T y \in [v_i, (1 + \epsilon)v_i]$, then clearly,

$$\sum_{i=1}^k u_i v_i \leq x^T (A + B) y \leq (1 + \epsilon)^2 \sum_{i=1}^k u_i v_i \quad (9)$$

For a fixed hyper-cube of the grid, consider the following LP based on the QP of (5)

$$\begin{aligned} \min: & e^T p + f^T q \\ \text{s.t. } & Ay \leq E^T p, \quad Fy = f, \quad y \geq 0 \\ & x^T B \leq q^T F; \quad Ex = e; \quad x \geq 0 \\ & u_i \leq x^T \alpha(i) \leq (1 + \epsilon)u_i; \quad v_i \leq \beta(i)^T y \leq (1 + \epsilon)v_i, \quad \forall i \end{aligned}$$

Algorithm. Run the above LP for each hyper-cube of the grid, and output an optimal point of the one giving the best approximation. As the number of hyper-cubes in the grid is $\text{poly}(\mathcal{L}, 1/\log(1 + \epsilon))$, the running time of the algorithm is $\text{poly}(M, N, \mathcal{L}, 1/\log(1 + \epsilon))$.

Correctness. Next we show that the above algorithm gives $(1 - \frac{1}{(1 + \epsilon)^2})$ -approximate NE of the game (A, B) . Let (x', y') be a NE of the given game, and (p', q') be such that $x'^T A y' = e^T p'$ and

$x'^T B y' = f^T q'$. Consider the hyper-cube containing $(x'^T \alpha(1), \dots, x'^T \alpha(k), \beta(1)^T y', \dots, \beta(k)^T y')$ of the grid and corresponding LP. Clearly, (x', y', p', q') is a feasible point of this LP and $\sum_{i=1}^k u_i v_i \leq e^T p' + f^T q' \leq (1 + \epsilon)^2 \sum_{i=1}^k u_i v_i$, since $e^T p' + f^T q' = x'^T (A + B) y'$. Therefore, at the optimal point $(\tilde{x}, \tilde{y}, \tilde{p}, \tilde{q})$ of the LP we get $e^T \tilde{p} + f^T \tilde{q} \leq (1 + \epsilon)^2 \sum_{i=1}^k u_i v_i$, and this gives,

$$\begin{aligned} \tilde{x}^T (A + B) \tilde{y} &\geq \sum_{i=1}^k u_i v_i \geq \frac{e^T \tilde{p} + f^T \tilde{q}}{(1 + \epsilon)^2} \quad (\text{using (9)}) \\ \Rightarrow e^T \tilde{p} + f^T \tilde{q} - \tilde{x}^T (A + B) \tilde{y} &\leq \left(1 - \frac{1}{(1 + \epsilon)^2}\right) (e^T \tilde{p} + f^T \tilde{q}) \end{aligned}$$

Let $\mu = \left(1 - \frac{1}{(1 + \epsilon)^2}\right)$, $\tilde{u} = \max_{\gamma \in X} \gamma^T A \tilde{y}$, and $\tilde{v} = \max_{\gamma \in Y} \tilde{x}^T B \gamma$. Clearly, $e^T \tilde{p} \geq \tilde{u}$ and $f^T \tilde{q} \geq \tilde{v}$ (using (1)). Let $D = e^T \tilde{p} + f^T \tilde{q} - \tilde{u} - \tilde{v}$, then $\tilde{u} + \tilde{v} - \tilde{x}^T (A + B) \tilde{y} = e^T \tilde{p} + f^T \tilde{q} - D - \tilde{x}^T (A + B) \tilde{y} \leq \mu(e^T \tilde{p} + f^T \tilde{q}) - D \leq \mu(e^T \tilde{p} + f^T \tilde{q} - D) = \mu(\tilde{u} + \tilde{v})$, since $\mu \in (0, 1)$. Therefore, (\tilde{x}, \tilde{y}) is a relative μ -approximate NE of the given game (A, B) (Definition 3).

Theorem 3. Let (A, B) be a rank- k game, and $A + B = \sum_{i=1}^k \alpha(i) \beta(i)^T$, such that $\alpha(i)$ s and $\beta(j)$ s are positive vectors. Then given an $\epsilon > 0$, a relative $\left(1 - \frac{1}{(1 + \epsilon)^2}\right)$ -approximate NE can be computed in time $\text{poly}(\mathcal{L}, 1/\log(1 + \epsilon))$, where \mathcal{L} is the input bit length. \square

For a symmetric game $(B = A^T, E = F, e = f)$, an (relative) ϵ -approximate symmetric NE can be defined as an (relative) ϵ -approximate NE with the same strategies, i.e., $x = y$. It is easy to check that, if we use the QP formulation of (6) instead of (5) in any of these algorithms, then the output strategy is an (relative) ϵ -approximate symmetric NE strategy.

5 Games with a Low Rank Matrix

In this section we show that if rank of even one payoff matrix (A or B) is constant, then Nash equilibrium computation can be done in polynomial time. Recall the best response polytopes P and Q (3) for the bilinear game (A, B) .

Lemma 6. Given a game (A, B) , there exists a vertex pair $((y, p), (x, q)) \in P \times Q$ such that (x, y) is a NE of (A, B) .

Proof. All the solutions of (4) over $P \times Q$ are the NE of the game, and existence of a solution is guaranteed (Proposition 1). Suppose $v \in P \times Q$ is a NE of (A, B) . It is easy to check that the entire face, formed by the set of tight equations at v , is solution, and it contains at least one vertex, as $P \times Q$ is bounded from one side. \square

Lemma 7. Let $k_1 = k_2 = k$ and $\text{rank}(A) = l$. The polytope P has at most $O(N^{l+k})$ vertices.

Proof. From (3), it is clear that P is in $(N + k)$ -dimensional Euclidean space, however $Fy = f$ gives k linearly independent equalities. Therefore, P is of dimension N , and at a vertex of P , N linearly independent inequalities must be tight. Since A is of rank l , $\text{rank}([A \ -E]) \leq l + k$. Therefore, $\exists! S \subset S_1$, $|S| = l + k$ such that $\forall i \in S_1 \setminus S, A_i y - p^T E^i \leq 0$ are not needed in defining the polytope P . At a vertex, if d inequalities are tight from S then rest $N - d$ must be of type $y_j = 0$, hence for a fixed $D \subset S$, $|D| = d$, there are at most $\binom{N}{N-d} = \binom{N}{d}$ choices to form a vertex. Therefore, the total number of vertices are at most $\sum_{i=0}^{l+k} \binom{l+k}{i} \binom{N}{i} \leq 2^{l+k} N^{l+k}$. \square

Note that if we remove the assumption $k_1 = k_2 = k$ then the exponent of N turns out to be a linear function of l, k_1 and k_2 . A similar proof can be worked out for Q .

Theorem 4. *If rank of either A or B is constant then a Nash equilibrium of a bilinear game (A, B) can be computed in polynomial time, assuming k to be a constant.*

Proof. Suppose $\text{rank}(A) = l$ (a constant) and $v = (y, p) \in P$ be a vertex. We can check in polynomial time whether v corresponds to a NE or not as follows. Let $S_x = \{i \in S_1 \mid A_i y - p^T E^i = 0\}$ and $S_y = \{j \in S_2 \mid y_j > 0\}$. Consider all $(x, q) \in \mathbb{R}^{M+k}$ such that

$$Ex = e; \quad \begin{array}{ll} \forall j \in S_y, & x^T B^j - q^T F^j = 0; \\ \forall j \notin S_y, & x^T B^j - q^T F^j \leq 0; \end{array} \quad \begin{array}{ll} \forall i \in S_x, & x_i \geq 0 \\ \forall i \notin S_x, & x_i = 0. \end{array}$$

Every such (x, q) lies in Q and makes a fully-labeled pair with v , and hence forms a NE (Lemma 1). Note that such an (x, q) can be obtained in polynomial time by solving an LP. Now the proof follows from Lemmas 6 and 7. A similar argument proves the other case when $\text{rank}(B)$ is constant. \square

As the set of bimatrix games is a subclass of the bilinear games (Example 1), where $k_1 = k_2 = 1$, Theorem 4 strengthens the results by Lipton, Markakis and Mehta [22] (Corollary 4), and Kannan and Theobald [17] (Theorem 3.2), where they require that the rank of both A and B to be constants. Note that k in Bayesian games depends on the number of types of players and in the sequence form, it depends on the number of information sets of players. Therefore, this result can be applied to these games if in their bilinear representation, a payoff matrix has low rank and k is constant.

In fact Theorem 4 gives a polynomial time algorithm to enumerate all the extreme equilibria of a bilinear game with a constant rank matrix, and an exponential time enumeration algorithm for any bilinear game. A similar (exponential time) algorithm was given by Avis et al. [2] to enumerate all Nash equilibria of a bimatrix game.

6 Conclusion

We have defined two-player *bilinear games*, where payoffs are represented by two matrices (A, B) and strategy sets are compact polytopes. In both bilinear and bimatrix games, the utilities are bilinear functions of strategy vectors. The scope of these games is large enough to capture many interesting classes of games like bimatrix games, two-player Bayesian games, polymatrix games, and two-player extensive-form games with perfect recall. Considering the rank-based hierarchy puts a structure on bilinear games, and by exploiting this structure and the similarity between bilinear and bimatrix games, we extended various combinatorial and algorithmic results, pertaining to the efficient computation of Nash equilibria, from bimatrix games to bilinear games. It will be interesting to know what other results of bimatrix games extend to bilinear games like 1) designing Lemke-Howson type algorithm for NE computation, 2) extending other algorithms for computation of approximate equilibria, etc.

Acknowledgments. We would like to thank Milind Sohoni for helpful comments and corrections.

References

1. Adsul, B., Garg, J., Mehta, R., Sohoni, M.: Rank-1 bimatrix games: A homeomorphism and a polynomial time algorithm. In: STOC. pp. 195–204 (2011)
2. Avis, D., Rosenberg, G.D., Savani, R., von Stengel, B.: Enumeration of Nash equilibria for two-player games. *Economic Theory* 42, 9–37 (2010)

3. Charnes, A.: Constrained games and linear programming. *Proceedings of the National Academy of Sciences of the USA* 39, 639–641 (1953)
4. Chen, X., Deng, X.: Settling the complexity of 2-player Nash-equilibrium. In: *FOCS* (2006)
5. Chen, X., Deng, X., Teng, S.H.: Computing Nash equilibria: Approximation and smoothed complexity. In: *FOCS*. pp. 603–612. IEEE Computer Society, Los Alamitos, CA, USA (2006)
6. Dantzig, G.B.: *Linear Programming and Extensions*. Princeton University Press (1963)
7. Daskalakis, C., Goldberg, P.W., Papadimitriou, C.H.: The complexity of computing a Nash equilibrium. In: *STOC*. pp. 71–78 (2006)
8. Daskalakis, C., Papadimitriou, C.: On a network generalization of the minmax theorem. In: *ICALP*, pp. 423–434 (2009)
9. Fudenberg, D., Tirole, J.: *Game Theory*. MIT Press (1991)
10. Glicksberg, I.L.: A further generalization of the Kakutani fixed point theorem, with application to Nash equilibrium points. *American Mathematical Society* 3(1), 170–174 (Feb 1952)
11. Goldberg, P.W., Papadimitriou, C.H.: Reducibility among equilibrium problems. In: *STOC*. pp. 61–70 (2006)
12. Goyal, V., Genc-Kaya, L., Ravi, R.: An fptas for minimizing the product of two non-negative linear cost functions. *Mathematical Programming* 126(2), 401–405 (2011)
13. Harsanyi, J.: Games with incomplete information played by “Bayesian” players, i-iii. part i. the basic model. *Management science* 14(3), 159–182 (1967)
14. Howson Jr, J.: Equilibria of polymatrix games. *Management Science* pp. 312–318 (1972)
15. Howson Jr, J., Rosenthal, R.: Bayesian equilibria of finite two-person games with incomplete information. *Management Science* pp. 313–315 (1974)
16. Immorlica, N., Kalai, A.T., Lucier, B., Moitra, A., Postlewaite, A., Tennenholtz, M.: Dueling algorithms. In: *STOC* (2011)
17. Kannan, R., Theobald, T.: Games of fixed rank: A hierarchy of bimatrix games. *Economic Theory* pp. 1–17 (2009)
18. Koller, D., Megiddo, N., von Stengel, B.: Fast algorithms for finding randomized strategies in game trees. In: *STOC*. pp. 750–759 (1994)
19. Koller, D., Megiddo, N., von Stengel, B.: Efficient computation of equilibria for extensive two-person games. *Games and economic behavior* 14, 247–259 (1996)
20. Kontogiannis, S., Spirakis, P.: Exploiting concavity in bimatrix games: new polynomially tractable subclasses. In: *APPROX*. pp. 312–325 (2010)
21. Lemke, C.E.: Bimatrix equilibrium points and mathematical programming. *Management Science* 11(7), 681–689 (May 1965)
22. Lipton, R., Markakis, E., Mehta, A.: Playing large games using simple strategies. In: *EC*. pp. 36–41. ACM New York, NY, USA (2003)
23. Nash, J.F.: Non-cooperative games. *The Annals of Mathematics* 54(2), 286–295 (1951)
24. Nisan, N., Roughgarden, T., Tardos, E., Vazirani, V. (eds.): *Algorithmic Game Theory*. Cambridge University Press, Cambridge, UK (2007)
25. Ponssard, J.P., Sorin, S.: The LP formulation of finite zero-sum games with incomplete information. *International Journal of Game Theory* 9, 99–105 (1980)
26. Shoham, Y., Leyton-Brown, K.: *Multiagent Systems: Algorithmic, Game-Theoretic, and Logical Foundations*. Cambridge University Press, New York (2009)
27. von Stengel, B.: Equilibrium computation for two-player games in strategic and extensive form. Chapter 3, *Algorithmic Game Theory*, eds. N. Nisan, T. Roughgarden, E. Tardos, and V. Vazirani pp. 53–78 (2007)
28. Vavasis, S.: Approximation algorithms for indefinite quadratic programming. *Mathematical Programming* 57, 279–311 (1992)